

AFRL-IF-RS-TR-2003-184
In-House Interim Report
July 2003



ADJUSTABLE BANDWIDTH CONCEPT (ABC) PERFORMANCE EVALUATION

Andrew J. Noga

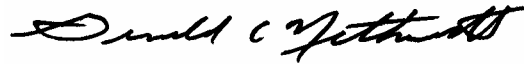
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

**AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK**

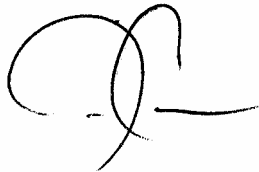
This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-2003-184 has been reviewed and is approved for publication.

APPROVED:



GERALD C. NETHERCOTT
Chief, Multi-Sensor Exploitation Branch
Info & Intel Exploitation Division
Information Directorate



FOR THE DIRECTOR:

JOSEPH CAMERA, Chief
Info & Intel Exploitation Division
Information Directorate

REPORT DOCUMENTATION PAGE			<i>Form Approved</i> OMB No. 074-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE JULY 2003	3. REPORT TYPE AND DATES COVERED In-House Interim, Oct 2001 – February 2003	
4. TITLE AND SUBTITLE ADJUSTABLE BANDWIDTH CONCEPT (ABC) PERFORMANCE EVALUATION			5. FUNDING NUMBERS PE - 62702F PR - 459E TA - H0 WU - C1	
6. AUTHOR(S) Dr. Andrew J. Noga				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) AFRL/IFEC 32 Brooks Road Rome, NY 13441-4114			8. PERFORMING ORGANIZATION REPORT NUMBER AFRL-IF-RS-TR-2003-184	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFRL/IFEC 32 Brooks Road Rome, NY 13441-4114			10. SPONSORING / MONITORING AGENCY REPORT NUMBER AFRL-IF-RS-TR-2003-184	
11. SUPPLEMENTARY NOTES AFRL Project Engineer: Andrew J. Noga/IFEC/315-330-2270/andrew.noga@rl.af.mil				
12a. DISTRIBUTION / AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.				12b. DISTRIBUTION CODE
13. ABSTRACT (Maximum 200 Words) Foundational work is presented regarding the objective performance analysis of the Adjustable Bandwidth Concept (ABC) Signal Energy Detector, U.S. Patent 5,257,211. The ABC pre-detection process is an effective method of combining time and frequency domain averaging to enhance the automated spectral survey task.				
14. SUBJECT TERMS Signal Detection, Automated Spectral Survey, Receiver Operating Characteristic (ROC), Cepstrum				15. NUMBER OF PAGES 44
				16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

Table of Contents

List of Acronyms	ii
1.0 Background and Motivation	1
1.1 Rationale for Problem Selection	2
1.2 Potential Impact	2
2.0 Approach	3
2.1 Uniqueness / Novelty of the ABC Approach.....	3
3.0 Existing Results	9
3.1 Directly Relevant Literature	9
3.2 False-Alarm Theoretical Considerations	10
3.3 Parameter Estimation for Achieving False-Alarm Rates.....	13
3.4 Signal Strength Estimation in the Presence of Noise	14
4.0 Research Results	15
4.1 Nonlinear Processing prior to Averaging	16
4.2 Test Scenario.....	18
4.3 Analytic Performance Evaluation Preliminaries.....	20
4.4 Analytic Performance Results	22
4.5 Experimental Performance Results.....	26
5.0 Conclusions.....	29
5.1 Future Work.....	30
5.1.1 Signal Grouping.....	30
5.1.2 Real-Time Implementation	30
5.1.3 Non-linear Compressors	30
5.1.4 Fusion of Detectors	31
5.1.5 SNR Estimation	31
5.1.6 Relationship to Other Processes	31
Appendix A: The Adjustable Bandwidth Concept (ABC) Algorithm	32
A.1 Averaging Over Frequency.....	33
A.2 Averaging Over Time	35
Appendix B: Filter Bank Representation of the ABC Process	36
Appendix C: Task Outline	38
References.....	39

List of Acronyms

ABC	Adjustable Bandwidth Concept
ADC	analog-to-digital converter
AF	Air Force
AFRL	Air Force Research Laboratory
AWGN	Additive White Gaussian Noise
BPSK	binary phase shift keyed
BW	bandwidth
CC	command and control
CF	center frequency
CITE	Center for Integrated Transmission and Exploitation
DFT	discrete Fourier transform
DSP	digital signal processing
EMC	electromagnetic compatibility
EMI	electromagnetic interference
FCC	Federal Communications Commission
FIR	finite impulse response
FFT	fast Fourier transform
IS	importance sampling
ISR	intelligence surveillance and reconnaissance
LPF	low-pass filter
ML	maximum likelihood
Pd	probability of detection
Pf	probability of false-alarm
PSD	power spectral density
RMS	root-mean-square
ROC	receiver operating characteristic
RST	rank-select-threshold
SNR	signal-to-noise ratio

1.0 Background and Motivation

In recent decades the Air Force, through the Air Force Research Laboratory (AFRL) Information Directorate (previously Rome Laboratory and Rome Air Development Center), has sponsored a substantial amount of research in the area of signal detection and parameter estimation. Likewise, industry and academia have invested significant efforts in these areas due to the benefits of such research to sonar, radar and communication applications. In particular, the Air Force has expended much effort to address the problem of automated, single-channel spectral survey. The goal of automated spectral survey is to process a digitized signal, typically from a step-tuned scanning receiver, to automatically (or semi-automatically) determine the presence of communications signals and their associated center frequencies (CFs), bandwidths (BW) and signal-to-noise ratios (SNRs). It is assumed that little or no a-priori knowledge exists regarding the digitized signal¹. The lack of a-priori knowledge gives rise to terms such as “blind signal detection” and “blind parameter estimation.” Within the scientific community as a whole, any digitized input signal can be of interest; other examples include speech and seismic signals. The present work is concerned with any simulated digital signal or analog signal acquired via an analog-to-digital converter (ADC). Through previous AF sponsorship, candidate software-based techniques for automated spectral survey were implemented for testing. However, no “acceptable” solution was found at that time; therefore, further development of these techniques was not performed. To be considered “acceptable” a technique must be computationally stable, require affordable computational resources, allow for ease of use by an operator, be timely, and be effective. Note that although spectral and parameter estimation techniques are abundant in the literature, automated blind signal *detection* remains relatively unaddressed.

Not surprisingly, problems with proposed solutions were associated with the activity detection function required for automated spectral survey. The processes that were developed utilized the log-spectrogram (a series of log-scaled periodograms) as pre-processing prior to detection and estimation of CF, BW and SNR for each component signal comprising the input signal. For the communications intercept application, the detection and estimation tasks were already being performed manually by an operator who would visually inspect such spectrograms or equivalent frequency panoramic displays on receive equipment. An appropriately trained operator can perform these tasks quite well in many signal scenarios. For these same scenarios, however, the automated algorithms tended to either break a single communication signal into multiple frequency bands, or to combine unrelated but spectrally close signals into single signals. An additional related problem had to do with the number of parameters associated with the software. Too many (and often confusing) software parameters left such techniques unattractive, even as an operator aid. Also, where spectrally overlapping, time-coincident signals were present, no means was provided in previously proposed techniques to help assess this co-channel scenario.

¹ Note that the digitized signal can actually have numerous signals-of-interest (or just portions of these signals) in the Nyquist band; the advancement of analog-to-digital converters allows for Megahertz-wide input bandwidths with sufficiently high dynamic range to collect many signals in a single scan step.

Based on knowledge of the above requirements and deficiencies of existing known techniques, the Adjustable Bandwidth Concept (ABC) signal energy detection technique was conceived, and ultimately patented (U.S. Patent 5,257,211). This technique is described in more detail in Appendix A and Section 2 and is the subject of this report.

1.1 Rationale for Problem Selection

Experimental tests and analyses of ABC algorithm implementations, first in the C language and then in Matlab (Mathworks, Inc.), demonstrated the effectiveness of the technique. Using simple threshold detection methods, visual inspection of a variety of input signal spectrograms and detection results led to the conclusion that these results corresponded quite well with human assessment of the spectrograms. Both computer-generated signal sequences and actual digitized signals were used during these experiments.

These subjective results were encouraging and provide motivation for further study to obtain quantitative performance measures. Although few parameters are required by the ABC algorithm, a more rigorous approach to the evaluation could also help to identify sensitivities to algorithm parameters. As is often the case, analytical results may also lead to algorithm modifications for enhanced performance. Thus the need was identified to quantify the detection results both with further experiments, and with analytical methods. In this report, detection probabilities and false-alarm rates for relevant ranges of SNR values have been established. These results are presented in the form of receiver operating characteristic (ROC) curves.

1.2 Potential Impact

The quantification of ABC algorithm performance as presented in this research serves to justify its usage in a variety of DoD/Air Force and commercial applications. It is a necessary step in the transition of this work to 6.2 (Exploratory Development). The primary Air Force applications are in Command and Control (CC) communication systems, and Intelligence Surveillance and Reconnaissance (ISR) systems. With regard to ISR, systems exist that fuse together a variety of information sources including ISR outputs to help situational awareness. These fusion systems require confidence measures for sensor outputs, particularly when dealing with ISR information. Without these confidence measures, the reliability or believability of the aggregate result remains in question. For CC, dissemination of information via communication links can be facilitated by channel assessments prior to transmissions. In both ISR and CC, intentional and unintentional interference occurs, particularly in dense signal environments. The ABC algorithm along with detection and parameter estimation can be used to either avoid contaminated channels, or to direct interference mitigation techniques. Other applications include automated testing for Electromagnetic Interference (EMI) and Electromagnetic Compatibility (EMC), TEMPEST testing and FCC compliance testing.

Given the recent growth in personal communication systems and the wide variety and number of electronic systems that emanate electromagnetic energy, “policing” the electromagnetic spectrum is a rather daunting task, and currently requires human assistance. Likewise, in both the government and commercial sectors, automated means of equipment compliance tests and evaluations can facilitate development and production. Furthermore, naturally occurring signals are often the subject of study with many aspects regarding conservation and the environment. Other natural signals such as the human voice have spectral properties that can be exploited. For example, bio-metrics taken from the human voice can aid in automated speaker identification. In fact, experiments on sampled human voice using the ABC algorithm, have demonstrated the utility of such a spectral characterization for quality assessment prior to speaker identification. Thus, a large number of applications exist for automated signal detection as provided by algorithms such as the ABC method.

2.0 Approach

There are two aspects to the proposed research, namely: i) motivation for selection of the ABC algorithm for pre-processing prior to signal detection, and ii) the needed quantitative performance analyses based on detection and false-alarm probabilities for appropriate input signal scenarios. The latter aspect is particularly important, and the tasks performed to accomplish this goal are covered within this report. In this section, both aspects of the proposed research will be addressed, but the emphasis will be on the general description of the ABC algorithm and its use for signal detection.

2.1 Uniqueness / Novelty of the ABC Approach

As alluded to in Section 1, the problem of automated signal detection and parameter estimation has been the subject of research sponsored by the Air Force over a number of years. A fundamental processing requirement for these resulting automated methods has been the generation of a time-frequency representation of the input signal (sequence). Due to its speed, well-understood properties and effectiveness, the short-term fast Fourier transform (FFT) is the most widely used transform to achieve a time-frequency representation of the input signal, particularly where real-time implementations are desired and little a-priori knowledge is available. However, other methods do exist. For example, the Chirp-Z transform has also been employed in some of the previously sponsored research. A separate identifiable area of relevant research is the investigation of these and many other proposed techniques for such representations, along with their application to automated signal detection and estimation of key signal parameters such as BW, CF and SNR. For the ISR and CC applications, experimental evidence suggests that the FFT-based approach yields highly satisfactory input signal representations for use by the human operator. In turn, the human operator can be effective at the signal detection and parameter estimation tasks in many scenarios. Therefore, the focus of this research is the automated interpretation of the resulting FFT-based time-frequency display.

In the previously sponsored research, proposed algorithms allowed for the averaging over time of power spectral estimates (from magnitude-squared FFT coefficients). This is analogous to the averaging capabilities of most spectrum analyzers for signals that *persist in time*. In effect, if spectral components remain relatively constant in amplitude at a given frequency for a period of time over which consecutive power spectral estimates are taken, then these estimates can be averaged together. Variations due to noise are thereby reduced, allowing for an enhanced probability of detection. Similarly, if the averaging over frequency bins of signals that are relatively short in duration can be accomplished, detection can be enhanced. Interestingly enough, even though this latter type of averaging is conceivable, this is not a typical capability in modern spectrum analyzers. A possible reason for this is the fact that *short duration signal events are easily missed by a human observer*. Moreover, the spectrum analyzer may not even respond to the event, particularly when the signal is aperiodic. Equipment developers have a clever solution to this problem which is effective in special cases. Because pulse-like signals can be repetitive (periodic) for many signal sources, the spectrum analyzer can be triggered by such events. This allows for the analysis of otherwise unobservable signals by time averaging. In contrast, automated signal detection can take advantage of averaging over frequency, when signals *persist in frequency*. Given that short duration signals will have wider bandwidths, averaging over frequency can enhance the probability of the detection of short duration signals. In terms of the probability of detection (P_d) and probability of false-alarm (P_f), a graph can be made relating the required input SNR versus number of estimates averaged, for a fixed P_d and some set of false-alarm probabilities. A notional version of this is shown in Figure 2-1.




Figure 2-1. Notional set of performance curves for “persistent” signals.

When arithmetically averaging power, actual graphs indicate that there is a “knee in the curve” somewhere between 10 and 20 estimates averaged, such that the benefit of averaging is reduced when exceeding this range. Averaging log-scaled power is also possible, and the utility of such averaging is addressed in the following results.

To gain a better appreciation for the benefits of averaging, some examples are presented. In the discussions of Figures 2-2a) through 2-3c), the upper portion of the figures show an input spectrogram (log scale) with frequency (bin) along the x-axis, and time (segment) along the y-axis. Each figure contains this spectrogram for comparison to detection results. In the lower portion of the figures are shown detections resulting from simple thresholding. Detections are shown in brighter shades, while dark regions indicate the absence of detections. The computer-generated input signal is known to consist of: a binary phase shift keyed (BPSK) component in the center of the band for all time segments; a stronger component of about 50 bins in width which turns on at about the 15th time segment and off at about the 48th time segment; a tone at around bin 330 for all time segments; a swept tone from bin 1 to bin 140; and 5 other tones of about 20 segments in duration that step across the band. Additive Gaussian noise is also present.

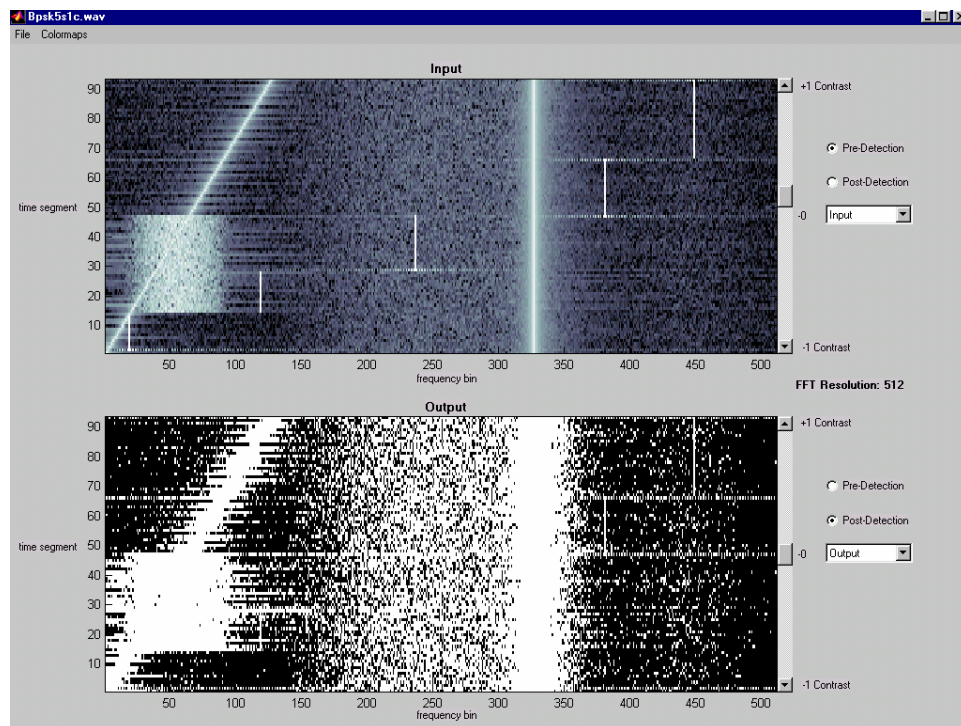


Figure 2-2a) Detection without averaging.

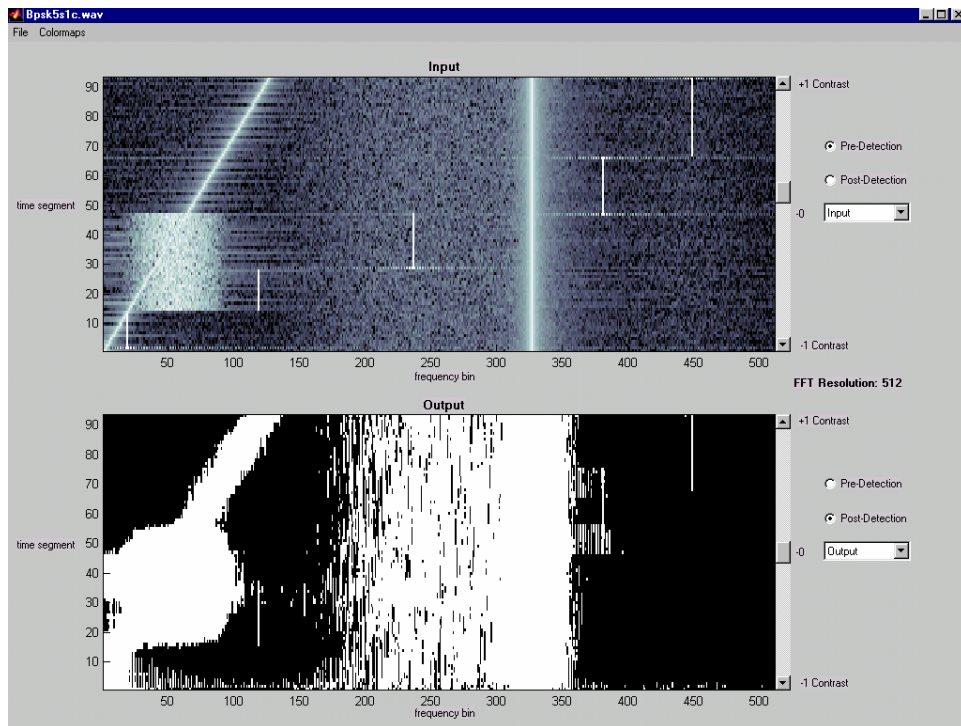


Figure 2-2b) Detection with time averaging only.

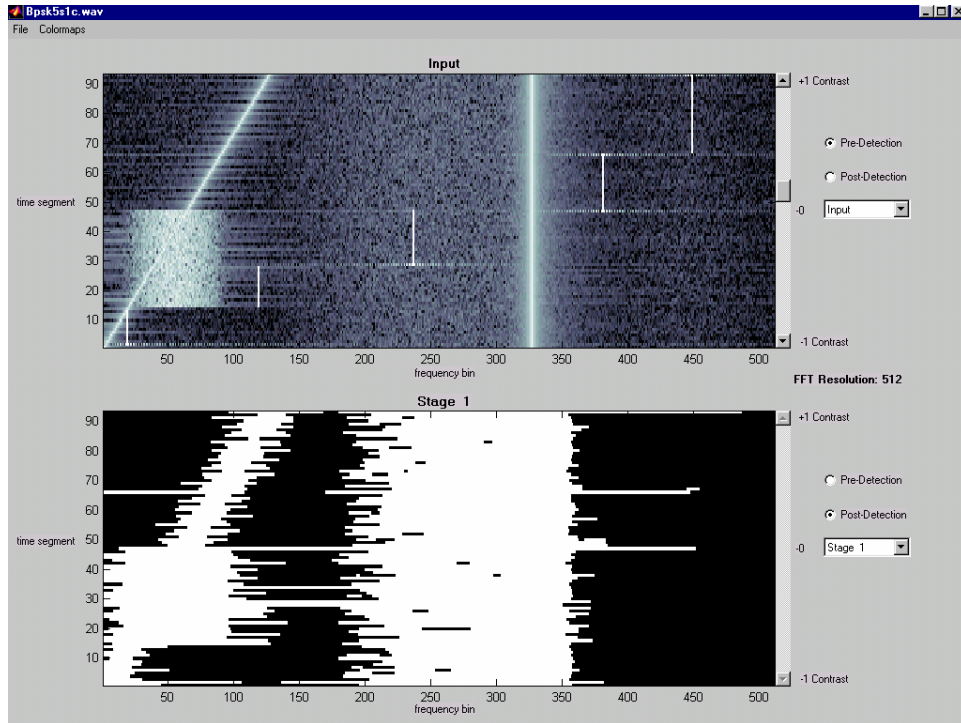


Figure 2-2c) Detection with frequency averaging only. (Also represents detection results on “wide bandwidth”, first stage ABC output.)

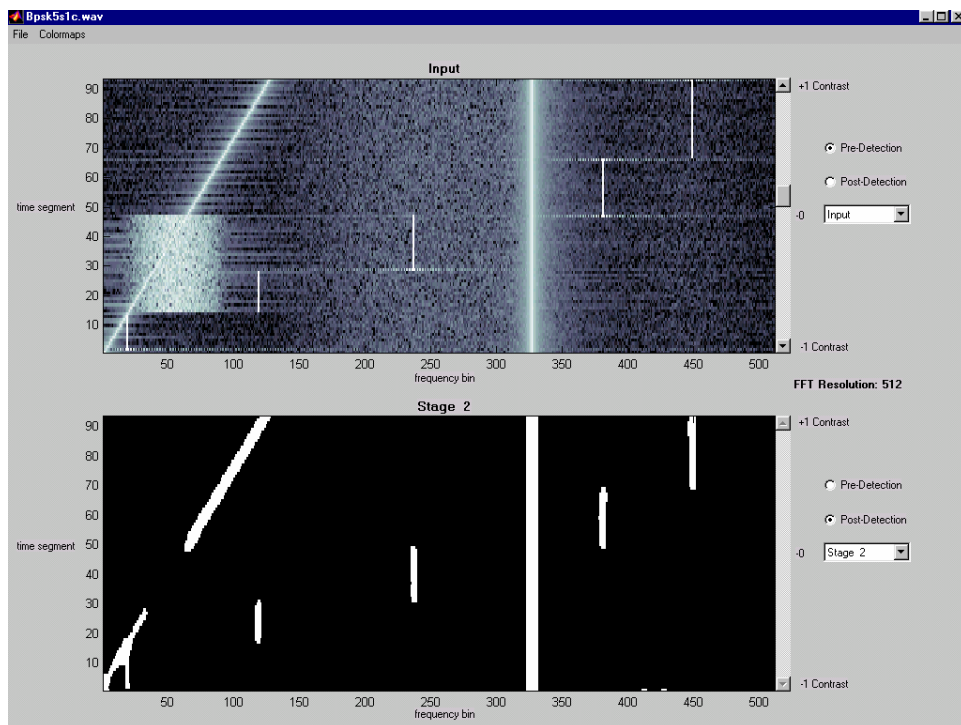


Figure 2-3a) Detection results on “medium bandwidth”, second stage ABC output.

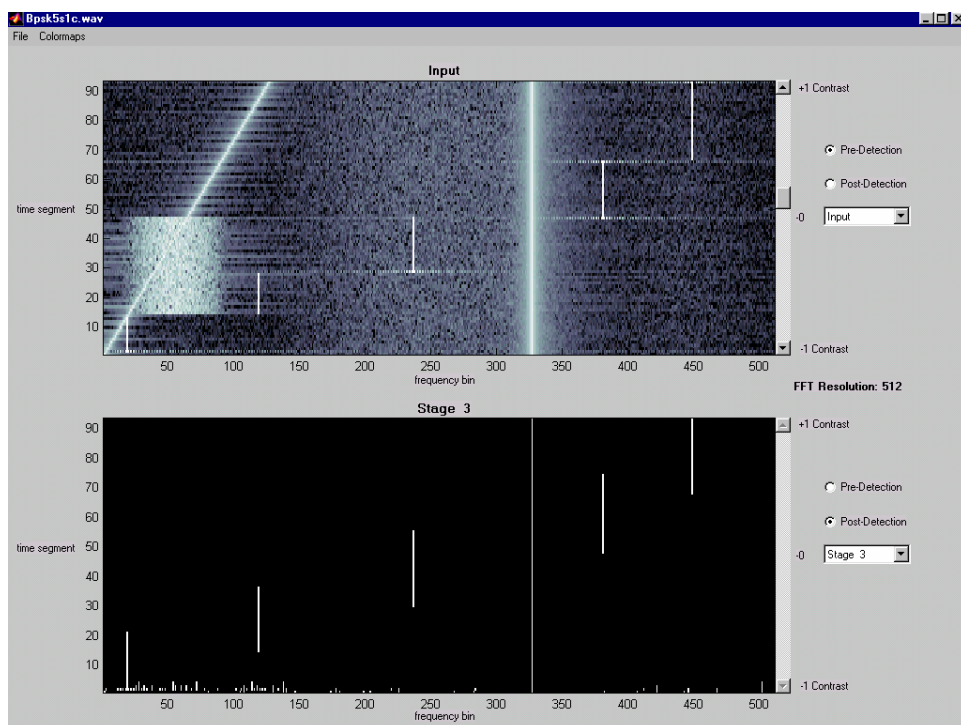


Figure 2-3b) Detection on “narrow bandwidth”, third stage ABC output.

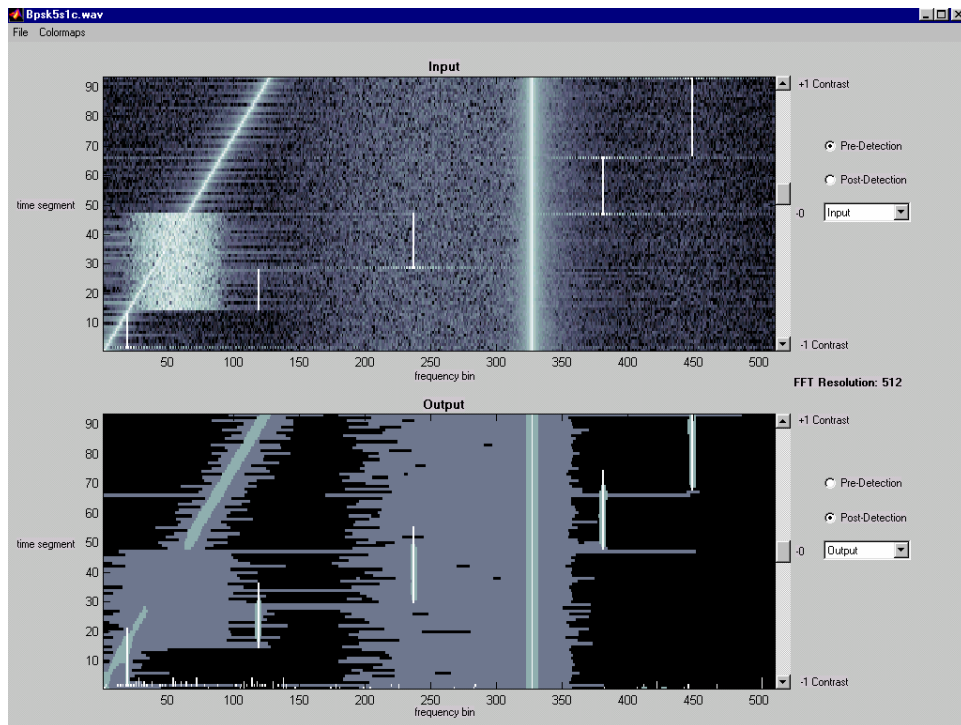


Figure 2-3c) Composite detection on all three ABC output stages.

Note that as seen in Figure 2-2a), detection results are difficult to interpret and in particular, the detected stepped tones are adjacent to numerous other detections and false-alarms. This makes it challenging, even with visual inspection by a human, to identify these tones. Also, the BPSK signal at the center of the band tends to be broken into random detections and misses. A trained human operator would be able to group these signal detections together, however, these detections would not be easily processed automatically. Finally, note that the tone located near bin 330 causes detections that range over 40 or more bins. This result may be “correct”, but important information is lost in the detection process. In fact, if the strength of the BPSK signal is increased, the tone detections are indistinguishable within the BPSK signal detections.

When time averaging is employed, detection results can be improved as shown in Figure 2-2b). False-alarms are reduced, and the detections caused by the BPSK signal at the center of the band are less random. However, short duration wide-band events associated with the start and stop of each stepped tone are no longer detected. Also, the detections caused by the tone near bin 330 have become indistinguishable from those caused by the BPSK signal.

When frequency averaging is employed, detection results can again be improved as shown in Figure 2-2c). False-alarms are reduced, and the detections caused by the BPSK signal at the center of the band are less random. However, narrow-band events associated with the stepped tones are no longer detected. Once again, the detections caused by the tone near bin 330 have become indistinguishable from those caused by the BPSK signal.

The advantages of ABC pre-processing prior to detection is graphically demonstrated in Figures 2-2c) through 2-3c). The detection was performed on a 3-stage ABC process, and results are shown for each stage. For stage 1, the threshold level is the same as that used for previous figures. Note that in this example, because no time averaging was included in the first stage of the ABC process, detection results are the same as when using frequency averaging only, as seen in Figure 2-2c). The detection threshold of any stage beyond stage 1 is set independently of all other stages. The data input to the detectors beyond stage 1 is residual data, and thresholds are therefore set in dB relative to zero. Results for these stages are shown in Figures 2-3a) and 2-3b). In general, as the stage number increases, less frequency averaging is performed, and more time averaging is performed. For the examples shown, stage 3 performs time averaging only. Thus, stage 1 is useful for detection of “wide band” signals, stage 2 is useful for detection of “medium bandwidth” signals, and stage 3 is useful for detection of “narrow bandwidth” signals. Figure 2-3c) shows all of the detection results overlayed on the same plot.

It can be seen that generally speaking, there are fewer false-alarms and detection has been enhanced. These observed results are typical for the ABC process and provide motivation for further quantified analyses. A Matlab implementation of the ABC process was used to generate the presented results (T. Hughes, AFRL-IF-RS-TM-1999-6, “A Signal Energy Detection Implementation,” December 1999). This implementation has been leveraged to provide experimental evidence for verification of analytical results obtained during the conduct of the present research.

3.0 Existing Results

In this section, background is given on relevant research already existing in the literature.

3.1 Directly Relevant Literature

Because the ABC technique operates on spectral amplitude data that has been logarithmically compressed, the outputs of the linear filter processes are related to the geometric mean. It is also important to note that as a pre-detection process, the ABC technique will be assessed based on false-alarm rates for a given probability of detection and input SNR. Particularly relevant publications include:

- [1] Rik Pintelon, et. al., “The Geometric Mean of Power (Amplitude) Spectra has a Much Smaller Bias than the Classical Arithmetic (RMS) Averaging,” *IEEE Transactions on Instrumentation and Measurement*, Vol. 37, No. 2, June 1988;
- [2] G. Corsini, et. al., “Cramer-Rao Bounds and Estimation of the Parameters of the Gumbel Distribution,” *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 31, No. 3, July 1995;

[3] Filippo Attivissimo, et. al., “A Study on Nonlinear Averagings to Perform the Characterization of Power Spectral Density Estimation Algorithms,” *IEEE Transactions on Instrumentation and Measurement*, Vol. 49, No. 5, October 2000.

Reference [1] discusses the statistical properties of the geometric average of power spectra compared to the more traditional root-mean-square (RMS) average. It represents an original proposal that the geometric mean has a much smaller bias than the RMS estimate, for the power spectra resulting from use of the discrete Fourier transform (DFT). The discussions in Reference [3] are supportive of the findings in [1], and also indicate that estimates can be produced that have a reduced variance by using the geometric average. Reference [2] discusses maximum likelihood (ML) estimates of the parameters of the Gumbel distribution which result at the output of logarithmic amplifiers used in some radar systems. Because the logarithmically compressed amplitude spectra from the output of the DFT are also Gumbel distributed when the input is noise-only, results presented in [2] also apply here. New findings in this report will be presented in the context of these and related existing findings.

3.2 False-Alarm Theoretical Considerations

The DFT- or FFT-based power spectral estimate, sometimes referred to as the periodogram, is considered to be a non-parametric approach. This is consistent with the notion of blind signal detection and estimation. However, some assumptions are appropriate for the purpose of developing a mathematically tractable analysis. Additive white Gaussian noise (AWGN) is assumed at the input, although the FFT process and the central limit theorem would allow us to loosen this requirement. Note that a false-alarm occurs when noise energy is declared to be signal energy. Therefore the false-alarm rate can be characterized by the distribution of the noise. In the noise-only case with a zero-mean Gaussian input of variance σ_t^2 , the magnitude-squared data from the periodogram bins are Exponentially distributed as

$$p_X(x) = \frac{1}{\mu_0} e^{-x/\mu_0}, \quad x > 0, \quad \mu_0 \propto \sigma_t^2. \quad (3-1)$$

The constant of proportionality between μ_0 and σ_t is determined by the FFT size and the window that is used. (See also [4].) The parameter μ_0 is assumed to be unknown. The nonlinear transformation, $Y = 10\log_{10}(X)$, results in Gumbel distributed data with

$$p_Y(y) = \frac{1}{\beta} \exp\left(\frac{y - \mu'_0}{\beta}\right) \exp\left(-\exp\left(\frac{y - \mu'_0}{\beta}\right)\right), \quad \mu'_0 = 10\log_{10}(\mu_0), \quad \beta = 10\log_{10}(e). \quad (3-2)$$

Expressions for the expected value of Y , $E(Y)$, and the variance of Y , $V(Y)$, are known; $V(Y)$ is a function of β only, which in our case, is constant! Restated, $V(Y)$ is not a

function of $\sigma_t > 0$. This is a rather interesting result! The magnitude squared and logarithmic transformations on the complex valued spectra transformed the variables from Gaussian distributions with a known mean of zero, to Gumbel distributions with a known scale parameter and thereby a known variance. Thus, it is now the mean of the resulting Gumbel distribution which gives us knowledge of the input noise power.

Numerical integration on the Exponential density function of (3-2) to examine the effects of logarithmic scaling serves to confirm the above results. Numerical estimates of both $E(Y)$ and $V(Y)$ are also consistent with theoretical results. Figure 3-1 shows example Exponential distributions and resulting Gumbel distributions. Note that the x-axis range for the Exponential distribution examples has increased by a factor of 10 when μ_0 increased by a factor of 10. The corresponding Gumbel distribution shows no change in variance. The increase in μ_0 translates to an increase in the location parameter to $10\log_{10}(\mu_0)$. Also note that the first “tail” in each of the curves representing the Gumbel distributions are at increased levels. This implies an increased probability for events in these ranges.

As it turns out, the Gumbel distribution has been known to the scientific community to be particularly useful for the modeling of important events occurring in nature. Because the Exponential distribution is a special case of the Weibull distribution, the Gumbel distribution is a member of the log-Weibull set of distributions, which are also referred to as Fisher-Tippet distributions and Extreme Value distributions. More specifically the Gumbel distribution is also called the Extreme Value Type I distribution. This latter name is more descriptive of what the distribution can model. It is used to describe various extreme condition phenomena such as floods, wind gusts and energy release during earthquakes. This appears to be an appropriate description of spectral events in the periodogram when the input is zero-mean Gaussian noise. Shown in Figure 3-2 is a log-scaled periodogram of zero-mean Gaussian data. Note the arrows pointing to some of the “extreme value” events which are predicted by the Gumbel distribution.

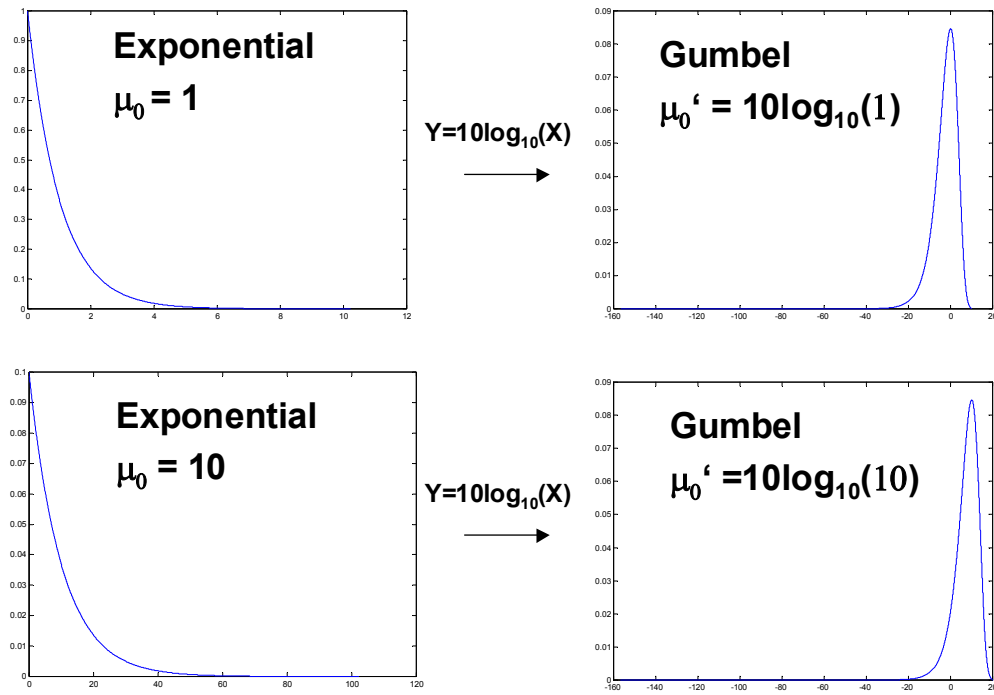


Figure 3-1. Examples of the effects of logarithmic compression on data distribution.




Figure 3-2. Results from a log-scaled periodogram.

3.3 Parameter Estimation for Achieving False-Alarm Rates

In order to automatically characterize the spectral estimate from the periodogram, receive-system calibration and/or parameter estimation will be required to achieve a specified false-alarm rate. With this in mind, the results of [2] apply. (The ML estimates for the parameters of the Gumbel distribution are also given in [5].) Note that interest in parameter estimation differs from performance evaluation using ROC curves. The latter is concerned with predicting or observing detection rates for a range of thresholds, and therefore for a range of false-alarm rates, at given SNR values. The former is the practical issue of estimating the noise power for the purpose of setting detection thresholds.

For the previously specified logarithmic transformation, the scale parameter estimate is $\hat{\beta} = \beta = 10 \log_{10}(e)$. However, in some applications the exact nature of the logarithmic compression may not be known. An example of this is the use of logarithmic amplifiers in radar systems. The ML estimates in this case become

$$\hat{\beta} = \frac{\sum_{i=1}^K y_i \exp\left(\frac{y_i}{\hat{\beta}}\right)}{\sum_{i=1}^K \exp\left(\frac{y_i}{\hat{\beta}}\right)} - \frac{1}{K} \sum_{i=1}^K y_i, \quad (3-3)$$

and

$$\hat{\mu}'_0 = \hat{\beta} \ln \left[\frac{1}{K} \sum_{i=1}^K \exp\left(\frac{y_i}{\hat{\beta}}\right) \right]. \quad (3-4)$$

Note that the K samples, y_i , can be from either consecutive times, consecutive frequencies, or both. The scale estimate in (3-3) must be solved in an iterative fashion, and used in (3-4) to then estimate the location parameter.

The above estimation technique is appropriate when the K samples are known to be from a noise-only input. In many practical scenarios, a signal component may exist in a substantial portion of the samples. In this case, the rank-select-threshold (RST) method of noise estimation can be used. In this method, the content of the bins from a set of periodograms are sorted in ascending strength, and the strength of the bin at around the p^{th} percentile is selected to estimate the noise. The fraction p is some appropriate percentage chosen based on the number of noise-only samples expected.

It should be noted that the estimate of Equation (3-4) is known to have a bias, but is asymptotically unbiased as K increases. An expression for this bias is given in [2] and results are shown in Figure 3-3 for values of K of interest in this report.




Figure 3-3. Bias on the estimate of the location parameter, $\hat{\mu}'_0$.

For $K = 1$, $E(Y) = \mu_0 + 0.57721\beta$ (the constant 0.57721 is Euler's number). Thus the bias in this case is $0.57721\beta = 2.5068$ dB.

3.4 Signal Strength Estimation in the Presence of Noise

Based on expressions given in Proakis [6], distributions resulting from the arithmetic averaging of spectral power can be given as

$$p_{\bar{X}/\bar{r}}(\bar{X}/\bar{r}) = \frac{K}{\mu_0} \left(\sqrt{\frac{\varphi}{\bar{r}}} \right)^{(K-1)} \exp(-K\varphi - K\bar{r}) I_{[K-1]}(2K\sqrt{\bar{r}\varphi}), \quad (3-5)$$

with $\varphi = \frac{\bar{X}}{\mu_0}$, and $\bar{r} = \frac{1}{K} \sum_{i=1}^K \frac{a_i^2}{\mu_0}$. Here, $I_{[K-1]}$ is the modified Bessel function of the first kind. The squared amplitude, a_i^2 , is the sum of the squared means of the real and imaginary components of the complex value obtained from bin i of the K selected FFT bins. The real and imaginary components are assumed to be Gaussian with variance $\sigma^2 = \mu_0/2$. Note that \bar{r} can be considered as the average signal-to-noise power ratio over the K independent samples. As should be expected, Equation (3-5) simplifies to Equation (3-

1) when $K = 1$ and $\bar{r} = 0$. The above expressions allow for the evaluation of the traditional RMS method of signal detection as will be described in Section 4.

In [1] the interest is to obtain an accurate estimate of the signal strength in the presence of noise when many thousands to a few million samples are available. It is recognized in the reference that when RMS averaging (i.e., arithmetic or linear averaging) is used to enhance the estimate, a bias is present that is much larger than when a non-linear averaging method is used. Specifically, the non-linear averaging technique used is the geometric mean resulting from the averaging of log-scaled power spectral magnitudes. Note that as previously presented for linear averaging, [2] provides an expression for the bias in the noise-only case such that the bias can be accounted for. Unfortunately, no such expression for the signal-plus-noise case is currently known to the author of this report. However, it can be inferred that such an expression will likely be a function of SNR, and will decrease as the number of estimates averaged, K , increases. Given that the bias is considered to be significantly larger when the arithmetic average is used, a particularly large bias will be present for the range of values for K of interest in this report.

4.0 Research Results

As more a-priori information regarding the input signal is available, this information can often be used to design a more mathematically optimal detection system. As a simple example, if a known waveform is to be detected in the presence of additive Gaussian noise, typically a matched filter or correlation receiver can be employed. In addition to being effective, such systems should: i) be readily implemented due to low complexity; ii) be computationally stable by being robust to small system parameter changes; iii) require affordable computational resources; iv) allow for ease of use by an operator and/or ease of integration into systems; and v) give timely results. Where less a-priori information is available, adaptive systems have been successfully employed, with trade-offs in the above desired system characteristics.

It's often the case that input signals of interest contain many signal components with little associated a-priori information. In fact, the acquisition process itself will result in a distorted representation of the original analog signal due to effects such as filtering and quantization. This leads to a substantial interest in non-parametric spectral estimation techniques. The FFT-based estimation technique is particularly attractive in the context of meeting the desired system characteristics.

Note that filtering (e.g., averaging of consecutive spectral estimates) can be very beneficial, but because of the lack of information regarding the input, it can also have unintentional and detrimental effects. The primary focus in this report is the detrimental effect filtering can have on the asynchronous detection of signal components. All time segments and frequency bins of the spectrogram are candidate locations for signal energy. Thus filtering over time or over frequency will reduce the resolvability of locations with noise energy in a neighborhood of locations near those of signal energy. As SNR increases, filtering will cause a significant number of locations in the neighborhood of a location containing signal energy to be falsely identified as containing

signal energy. Depending on the filter impulse response, a de-sensitization can also occur. To address these concerns, amplitude compression is considered prior to filtering.

4.1 Nonlinear Processing prior to Averaging

Under the premise that filtering in the presence of nearby signals becomes a mechanism for misclassification, the system of Figure 4-1 is considered. Of interest are those functions, g , which operate on the magnitude-squared spectral values, X , such that averaging by the linear filter will result in enhanced detection performance. In the figure, k_i represents the k^{th} frequency bin of the i^{th} time segment. The output of the threshold detector is the binary detection grid, $B(k_i)$, which contains a 1 in locations where signal energy is detected and a 0 otherwise.




Figure 4-1. An automated blind signal detection system.

The following interrelated properties of a continuous compression function g can be identified:

Property 1. $g(x)$ must be one-to-one, and onto as a map from the domain $x > 0$ to the range $x > 0$ or $-\infty < x < +\infty$.

Property 2. $g(x)$ must be strictly increasing in the domain $x > 0$. This implies that g is one-to-one and has an inverse, and implies that $\frac{dg}{dx} > 0$ for $x > 0$.

Property 3. $g(x)$ must be concave down in the domain $x > 0$. This implies that $\frac{d^2g}{dx^2} < 0$ for $x > 0$.

Property 4. $g(x)$ minimizes the positive relative difference, $\frac{g(x_b) - g(x_a)}{g(x_b)}$, $0 < x_a < x_b$.

These properties are required for a continuous compression function that minimizes misclassifications that occur due to filter transients.

Properties 1 through 3 motivate the consideration of functions such as $g_p(x) = x^{1/p}$, where $p > 1$. From Property 4, the value of p is sought which minimizes $1 - (x_a / x_b)^{1/p}$. The solution is $p = +\infty$, however, this leads to the trivial case where $g(x) = \lim_{p \rightarrow \infty} x^{1/p} = 0$. This result does not satisfy the remaining properties, 1 through 3.

In Section 3 it has been observed that logarithmic compression prior to averaging produces useful results. This motivates the consideration of the function $g(x) = \ln(x)$, the natural log of x . It is easily shown that this function satisfies Properties 1 through 3. To make relevant observations regarding Property 4, it is recalled that

$$e^x = \lim_{p \rightarrow \infty} \left(1 + \frac{x}{p}\right)^p. \quad (4-1)$$

Manipulation of Equation (4-1) results in

$$\ln(x) = \lim_{p \rightarrow \infty} p \cdot (x^{1/p} - 1), \quad (4-2)$$

which encourages the consideration of $g_p(x) = p \cdot (x^{1/p} - 1)$. The quantity

$$\frac{g_p(x_b) - g_p(x_a)}{g_p(x_b)} = 1 - \frac{x_a^{1/p} - 1}{x_b^{1/p} - 1}$$

is minimized as $p \rightarrow +\infty$. This shows that for this set of functions, $g_p(x) = p \cdot (x^{1/p} - 1)$, the function $g(x) = \ln(x) = \lim_{p \rightarrow \infty} g_p(x) = \lim_{p \rightarrow \infty} p \cdot (x^{1/p} - 1)$ is a function from the set that minimizes the positive relative difference of Property 4.

Well known properties of the logarithm can be applied to further show that bases other than base e can be chosen. Note that nothing has been assumed regarding the distribution of the positive real random variable X in the derivation of the above result. This does not imply, however, that all signal scenarios will result in the conclusion that logarithmic compression should be used. Analytical and experimental analyses that follow will yield further insight into performance issues. The special case to be considered is the moving average filter, which in combination with logarithmic pre-compression, results in the geometric average. To be more general, functions from the set

$$g_p(x) = f \cdot (x^{1/p} - c), \quad p > 1 \quad (4-3)$$

can be considered, where f is some scaling factor that is a function of p only, and c is a constant. From this set, aspects of both logarithmic compression and square-root compression ($f=1, c=0, p=2$) will be considered and compared to the RMS average.

4.2 Test Scenario

The determination of an appropriate test scenario for the performance evaluation of automated blind signal detection processes is a non-trivial task. The difficulty in defining a test scenario is largely due to the fact that little is assumed known about the input sequence. However, there are reasonable assumptions that can be made which both allow for a meaningful performance evaluation and lead to mathematically tractable analyses. As previously indicated, the input is assumed to consist of signals in additive white Gaussian noise. Because the detection processes considered herein are discrete-time processes, the term “white” refers to the fact that the FFT will produce frequency bins that contain noise components with uniform power over frequency (with the exception of the first bin and a bin near the sampling frequency for real inputs). Thus the complex spectral coefficients are also Gaussian and independent, leading to the density function for spectral power given in Equation (3-5).

With the statistical nature of the spectrogram established, there remains a need to select the deterministic signal components of the spectrogram to be used for performance testing. A logical candidate is the delta function (in this case an impulse in the spectrogram) which is particularly effective at evaluating filter responses. Because averaging over time segments and over frequency bins is a simple filtration process, an impulse in varying levels of noise could be considered. This could potentially help in the assessment of misclassification due to filter transients. However, it can also be noted that the matched filter processing of such a spectrogram would be one in which no averaging is employed. This would render all pre-detection processes which employ averaging, as sub-optimal techniques. This is contrary to the reason for averaging in pre-detection processes, namely, to enhance the probability of detection of persistent signals. Thus another fundamental assumption for the research herein is that persistence in time, frequency or both is a predominant characteristic of the unknown input signal components.

Based on the above considerations, the proposed deterministic component of the test spectrogram is that of a uniform amplitude rectangular pulse with duration K samples. This pulse can persist for K time segments or K frequency segments, but at least K samples containing noise-only exist before and after the samples containing the pulse. Whether the pulse is over time or over frequency will not be relevant in the results presented; therefore the pulse will more generally be referred to as being persistent for K consecutive samples. It is important to note that the locations of such pulses within the spectrogram would be unknown. The test scenario is therefore considered to be an asynchronous spectral pulse. However, when comparing experiments with analytic results, there is no need to randomize the locations of simulated pulses if detections and false alarms are properly accounted for. In order to make detection performance assessments, the concepts of detections and false alarms must first be better defined.

If the power values from the spectrogram are not averaged prior to threshold comparison and detection classification, determining detection and false alarm rates is straightforward. A frequency bin containing signal-plus-noise is defined to be class “1”, and a bin containing noise-only is defined to be class “0”. Declaring a “1” when the bin

is defined to be from class “1” is a correct-detection (or simply referred to as a detection), and declaring a “1” when the bin is defined to be from class “0” is a false-alarm. The difficulty in defining classes arises when filtering is employed prior to detection. Filtering does not directly reduce the number of samples for which a class determination is required. (In fact, the ABC pre-detection process increases the number of threshold comparisons needed by a factor of M , the number of stages. See also Appendix A.) Rather, filtering has the effect of delaying the output relative to the input. If an analysis window of K consecutive samples of a filtered spectrogram is considered, there is a correspondence to the pre-filtered spectrogram that allows for classes to be defined. The sample delay, d , of the moving average filters considered herein is found to be $d = (K - 1)/2$. More generally, any symmetric finite impulse response (FIR) filter of order $K-1$ (i.e., with K coefficients) has a sample delay of $d = (K-1)/2$. By taking into account this sample delay, classes for the samples at the filter output can be defined based on the classes of the filter input samples. Table 4-1 summarizes these class definitions as an analysis window of size K is slid over a pulse that persists for K samples.

Case	Window Samples Containing Signal-Plus-Noise	Window Samples Containing Noise-Only	Class Definition
1	K (odd)	0	1
2	$K-1$	1	1
3	$K-2$	2	1
⋮	⋮	⋮	⋮
⋮	⋮	⋮	⋮
⋮	⋮	⋮	⋮
$d+1$	$d+1$	$d=(K-1)/2$	1
$d+2$	d	$d+1$	0
$d+3$	$d-1$	$d+2$	0
⋮	⋮	⋮	⋮
⋮	⋮	⋮	⋮
⋮	⋮	⋮	⋮
K	1	$K-1$	0
$K+1$	0	K	0

Table 4-1. Class definitions for a moving average filter of order $K-1$.

Table 4-1 enumerates class definitions for each sample output generated as the average is moved over input samples. The first row of the table is the case where the analysis window is aligned with the pulse and represents the matched filter case. The rows progress with fewer signal-plus-noise bins and more noise-only bins being averaged together. This process is based on the “shift-multiply-add” description of discrete-time convolution. The last row also represents the matched filter case, but for the noise-only input. Note that by considering values of K that are odd, the ambiguous case is avoided

where an equal number of signal-plus-noise and noise-only bins are averaged together. It also allows for the cases to be paired for detection and false-alarm calculations. For example, case 1 can be paired with case $K+1$ to represent ideal (matched filter) performance. Likewise, case 2 can be paired with case K , case 3 can be paired with case $K-1$, and so on. These latter pairings represent the adverse effects of filtering where performance is degraded relative to the matched filter. Table 4-2 shows the list of pairs that are considered. By averaging together the performance results of the ensemble of pairs, an expected degraded performance is obtained for a given value of K . Alternatively, the density functions representing class “0” can be averaged together for overall false-alarm rate calculation, and the density functions representing class “1” can be averaged together for overall detection rate calculation. Thus each of the cases are considered to be equally likely to occur, consistent with the asynchronous nature of the spectral pulse.

Case Pairs
$\{1, K+1\}$
$\{2, K\}$
$\{3, K-1\}$
\vdots
\vdots
\vdots
$\{d+1, d+2\}$

Table 4-2. Case pairings for ROC curve generation.

4.3 Analytic Performance Evaluation Preliminaries

To facilitate the generation of ROC curves for performance evaluations and comparisons, the density function in Equation (3-5) is transformed using $\hat{\bar{x}} = \ln(\bar{x})$. The resulting density function is

$$p_{\hat{\bar{x}}/\bar{r}}(\hat{\bar{x}}/\bar{r}) = K\gamma \left(\sqrt{\frac{\gamma}{\bar{r}}} \right)^{(K-1)} \exp(-K\bar{r}) \left[\exp(-K\gamma) I_{[K-1]}(2K\sqrt{\bar{r}\gamma}) \right], \quad (4-4)$$

with $\gamma = \frac{\exp(\hat{\bar{x}})}{\mu_0}$, and $\bar{r} = \frac{1}{K} \sum_{i=1}^K \frac{a_i^2}{\mu_0}$. Note that ROC curves for the arithmetic averaging of magnitude-squared spectral values can be generated using either Equation (3-5) or Equation (4-4). However, the term $I_{[K-1]}$ in the equations can cause severe computational instabilities when numerical integration is used to evaluate detection and false-alarm

probabilities. To address this problem, the terms bracketed in Equation (4-4) are converted to an alternate form.

An integral form of the modified Bessel function of the first kind is

$$I_{[K-1]}(\eta) = \frac{1}{\pi} \int_0^\pi \exp(\eta \cos(\theta)) \cos([K-1]\theta) d\theta, \quad (4-5)$$

when $K-1$ is integer valued. Because K is integer valued herein, an expression for the bracketed terms in Equation (4-4) is

$$\exp(-K\gamma) I_{[K-1]}(2K\sqrt{\bar{r}}\gamma) = \frac{1}{\pi} \int_0^\pi \exp(2K\sqrt{\bar{r}}\gamma \cos(\theta) - K\gamma) \cos([K-1]\theta) d\theta. \quad (4-6)$$

By using Equations (4-4) and (4-6) to numerically evaluate the relevant density functions, computational stability is achieved without the need to resort to more involved non-analytical techniques such as Importance Sampling (IS). For the results herein, values of interest are $K \leq 20$ and $\bar{r} \leq 10$. Note that in Equation (3-5) and therefore in Equation (4-4), the parameter $\mu_0 > 0$ can be set to 1 for convenience. A step size of 1e-3 over the range $-40 \leq \hat{x} \leq 10$ and rectangular integration were employed for the integral of Equation (4-4). With a step size of 1e-2 and by employing trapezoidal integration for the integral of Equation (4-6), overall numerical results were found to be quite accurate. The method chosen to test accuracy was to check the numerical integration of the density functions over the range using rectangular integration. Because the interest is in false alarm rates at about 1 per thousand, it was verified that the density functions integrate to 1 with at least 6 decimal places of accuracy. In fact, the results were typically found to be greater than 10 places of accuracy. Thus the combined approach of first converting to natural-log scale and recognizing the relationship in Equation (4-6) becomes a rather unique and effective way of overcoming computational accuracy problems for the required numerical integrations.

To generate the ROC curves for the geometric averaging of magnitude-squared spectral values, Equations (4-4) and (4-6) can be used to derive the appropriate density functions. The method used herein involved first setting $K = 1$ in the equations, and subsequently employing numerical convolutions. Two cases can be identified, one with an arbitrary parameter $0 < r_i \leq 10$ representing the signal present case, and one with $r_i = 0$ representing the noise-only case. For both cases, the relevant density function becomes

$$p_{\hat{X}/r_i}(\hat{X}/r_i) = \gamma \exp(-r_i) [\exp(-\gamma) I_{[K-1]}(2\sqrt{r_i}\gamma)], \quad (4-7)$$

with $\gamma = \frac{\exp(\hat{x})}{\mu_0}$, and $r_i = \frac{a_i^2}{\mu_0}$. The sum of K independent samples results in

$$p_{\hat{X}/\bar{r}}(\hat{X}/\bar{r}) = p_{\hat{X}/r_1}(\hat{X}/r_1) * p_{\hat{X}/r_2}(\hat{X}/r_2) * p_{\hat{X}/r_3}(\hat{X}/r_3) \dots * p_{\hat{X}/r_K}(\hat{X}/r_K), \quad (4-8)$$

with $\bar{r} = \{r_1, r_2, r_3, \dots, r_K\}$. For the specific scenario where uniform weighting is used, r_i is either zero (noise-only) or some constant SNR. Note that because of the dependence of the modified Bessel function, $I_{[K-1]}$, on the ratio r_i , the convolution process does not allow for expression simplification by the use of the average, \bar{r} , as a parameter. Also note that scaling Equation (4-8) by the factor $1/K$ to produce the average rather than the sum, will have no affect on the ROC curves subsequently generated.

4.4 Analytic Performance Results

The input test scenario of Section 4.2 is now used for performance comparisons. A detector employing arithmetic averaging of power resulting in the density function of Equation (4-4), is compared to a detector employing geometric averaging resulting in the density function obtained from Equation (4-4) with $K = 1$ and the convolution process of Equation (4-8). Results are shown in Figures 4-2a) through 4-4b). Note that a false alarm probability that is commensurate with the probability of a miss (1 minus the probability of a detection) will be referred to as a “large” false alarm probability. In demodulation scenarios false alarm and miss errors are both considered as bit errors. Detection thresholds are set to achieve an equal number of false alarms and misses. Thus in a relative sense, demodulation performance is measured at large false alarm probabilities. This is in contrast to radar, sonar and automated signal detection processes such as the ABC algorithm, where the probability of a false alarm must be small in an absolute sense. For these processes, false alarm probabilities on the order of 1 or less per thousand may be of interest.

An SNR per bin (sample) of 1 is first considered. As seen in Figures 4-2 a) and b), there is a slight but distinct performance advantage in using arithmetic averaging relative to geometric averaging at large false alarm probabilities. The figures also show the degraded performance expected because of transients associated with the averaging filter. Arithmetic averaging under degraded conditions still outperforms geometric averaging under degraded conditions, at these large false alarm probabilities.




Figure 4-2a). Pulses persistent for $K = 3$ consecutive samples at SNR = 1, large Pf range.




Figure 4-2b). Pulses persistent for $K = 7$ consecutive samples at SNR = 1, large Pf range.




Figure 4-3a). Pulses persistent for $K = 3$ consecutive samples at SNR = 1, small Pf range.




Figure 4-3b). Pulses persistent for $K = 7$ consecutive samples at SNR = 1, small Pf range.




Figure 4-4a). Pulses persistent for $K = 3$ consecutive samples at SNR = 7, large Pf range.




Figure 4-4b). Pulses persistent for $K = 3$ consecutive samples at SNR = 7, small Pf range.

However, when considering small false alarm probabilities as seen in Figures 4-3a) and 4-3b), results are significantly different. It can be seen that both arithmetic averaging and geometric averaging suffer from degraded performance due to filter transients. In this scenario, degraded geometric performance is better than the degraded arithmetic performance. It is also apparent that as K is increased from 3 to 7, under degraded conditions, geometric averaging continues to outperform arithmetic averaging.

For reference, it is noted that without averaging, the detection probability is about .02 for a false alarm probability of $1e-3$ at this SNR. Based on the expected degraded performance shown in the figures, this would seem to imply that there is no advantage to either type of averaging. The expected degraded performance is about .02 to .025 in the given scenario. What is not readily apparent in the figures is the fact that the ideal performance will be achieved at some sample. Ideal performance is defined herein as the matched filter case, described in Section 4.2. Also, near ideal performance will be achieved for a significant number of samples, when K is greater than 3.

In Figures 4-4a) and b), the SNR per bin has been increased to 7 and K is 3. As seen in Figure 4-4a), the ideal detection probability is essentially 1 for both arithmetic and geometric averaging, for this scale of false alarm probabilities. Note that for false alarm probabilities below 0.2, the degraded performance of the geometric average is better than the degraded performance of the arithmetic average. For small false alarm probabilities as shown in Figure 4-4b), results are more dramatic. Degraded geometric performance is nearly 100 percent better than the degraded arithmetic performance. It is expected that as SNR increases, transients due to average filtering will be more pronounced. Geometric averaging then becomes particularly beneficial.

Once again for reference, it is noted that without averaging, the detection probability is about 0.57 for a false alarm probability of $1e-3$ at $SNR = 7$. This is somewhat better than the degraded geometric average performance of about 0.49. However, it is again noted that ideal performance will be achieved at some sample. Based on the figure, that detection performance is about 0.98.

4.5 Experimental Performance Results

To help confirm previously presented analytic results, experiments have been conducted. Experiments consisted of Monte-Carlo computer simulations of signals of interest and detection results. Through simulations, insight is also gained into an alternative detection process, where square-root compression is employed in the process of Figure 4-1. Representative results are given.

Shown in Figure 4-5 is a comparison of experimental and theoretical performance results, for $K = 7$ and an SNR of 1. Due to the length of time associated with the simulations, comparisons are limited to large values of probability of false alarms. Note that experiments are in agreement with theory. Several other selected scenarios also served to confirm the analytic results.




Figure 4-5. Pulses persistent for $K = 7$ consecutive samples at SNR = 1.0, large Pf range.




Figure 4-6. Pulses persistent for $K = 7$ consecutive samples at SNR = 10, large Pf range, degraded performances at a sample delay of 3.

The analyses presented in Section 4.1 also served as motivation to consider employing the square-root as a compression function. Simulations were easily modified to include square-root compression of the magnitude-squared spectral values. The experimental ROC result for a selected scenario is shown in Figure 4-6. The results presented correspond to case 4 of Table 4-1, where 4 samples containing signal-plus-noise were averaged with 3 samples containing noise-only (and $K = 7$). An SNR of 10 was selected for presentation. The significant observation is the fact that square-root compression before averaging, also referred to in the figure as the “root-arithmetic” average, performs much better than both the arithmetic and geometric averages for this scenario. This is consistent with the fact that most envelope demodulation schemes in practice use the square-root of magnitude-squared data, prior to filtering.

In Figures 4-7 a) and b) is shown a comparison between experimental and theoretical performance results for the previous example of Figure 4-4a), where an SNR per sample of 7 is considered, and $K = 3$. Note the close agreement between theory and experiment for the arithmetic and geometric averages. Also included is the experimental result for root arithmetic averaging. Again it can be seen that the root-arithmetic average performs much better than both the arithmetic and geometric averages for this scenario, when demodulation is of interest. However, initial experimentation of the ABC pre-detection process with root-arithmetic averaging employed did not yield particularly useful results. There did appear to be an advantage over arithmetic averaging, but visual comparisons of the input spectrogram and the detection results indicated that geometric averaging gave much better results. This is an area that will require further research.




Figure 4-7a). Pulses persistent for $K = 3$ consecutive samples at SNR = 7, large Pf range; experimental and theoretical results for expected degraded performances.




Figure 4-7a). Pulses persistent for $K = 3$ consecutive samples at $\text{SNR} = 7$, large P_f range; experimental and theoretical results for expected degraded performances. (Zoom view.)

5.0 Conclusions

The Adjustable Bandwidth Concept (ABC) technique (U.S. patent 5,257,211), has been presented as a pre-detection process for discrete-time signals. The technique has previously been found to be quite useful for the off-line processing of collected signals, including audio, communications and radar signals. The work herein has been necessary to better understand the quantitative performance aspects of the ABC process before further use and before committing to hardware-based real-time implementations.

The intended application of the ABC process is to allow for both averaging over time and over frequency, when signal components persist in these domains. Generalizations to other domains such as the spatial domain may be possible, however, this was not investigated in this research. Little or no a-priori knowledge is assumed regarding the input signal and its components.

Toward the goal of objective performance evaluation of the process, receiver operating characteristic (ROC) curves have been generated for the general class of automated spectral analysis tools that employ the FFT and logarithmic compression of the magnitude. Other compressions have been considered, however, mathematical evidence given herein along with recent literature suggests that logarithmic compression results in particularly good performance.

In developing a methodology for performance evaluation, some key results have been obtained that may prove useful in a wide variety of signal detection systems. First,

a test scenario has been identified which allows for the fair comparison of competing detection methods. This scenario consists of a set of uniform pulses that occur in a random fashion in the test signal, such that the location of the pulse (time or frequency location) is unknown. The asynchronous nature of the pulses forces the consideration of an important problem that arises in blind signal detection systems, namely, misclassifications due to filtering. While it is true that filtering is beneficial, if not done properly, filtering results in pulse spreading leading to false-alarms and misses that occur in the neighborhood of the pulse.

A second key methodology result is a formulation of the statistical properties of the spectral data, that leads to a computationally stable numerical calculation of false alarms and detections. This avoids the need for time consuming simulations that often occur when Monte-Carlo analyses are employed.

5.1 Future Work

In this sub-section, some ideas are presented for future work that can be done to both further the application of the ABC process, and to gain additional insight into performance potential.

5.1.1 Signal Grouping

As described, the ABC algorithm is a pre-processing step prior to detection. However, even with great detection results, an interpretive step is still required. For example, the various detections that result are obtained on a bin-by-bin and segment-by-segment basis. Methods of combining adjacent detections, and possibly even non-adjacent detections, are required to automatically identify signal components within the log-spectrogram (see also [7]). It may be that magnitude and phase information, along with some a-priori knowledge of input signal structures can help in this regard. Another possibility is that image processing concepts/techniques may apply, given that visual inspection by a human can result in good detection combinations.

5.1.2 Real-Time Implementation

Favorable off-line experimental results have caused an interest in creating real-time implementations. This can expedite experimental performance evaluations, and allow for a wider variety of signal tests. The main computational components of the ABC process are the FFT, conversion to magnitude, and logarithmic compression. All computational components are readily implemented or are available in commercial DSP products.

5.1.3 Non-linear Compressors

Although this report has focused on the logarithmic compressor prior to averaging, other compression functions are conceivable and potentially useful. Initial experimental results regarding the use of square-root compression indicated a poorer performance relative to the logarithmic compression. However, ROC curves should be generated to confirm these results, particularly at low false-alarm probabilities. Also, for real-time implementations, it may not be possible to achieve exact logarithmic compression at desired speeds. This may lead to consideration of piece-wise differentiable functions that approximate the performance of the logarithmic compressor.

5.1.4 Fusion of Detectors

Recent literature suggests that the fusion of a set of detection processes that operate on the same data, can potentially enhance ROC performance [8]. There may be merit in considering such a technique with the ABC process. In particular, useful compressors can be chosen to generate a parallel set of ABC processes, each with a unique compressor. Furthermore, it may be beneficial to have a parallel set of ABC processes, each with a specific set of system parameters to allow for various input scenarios. Efforts would need to focus on identifying useful ABC configurations and methods of combining (fusing) detection results.

5.1.5 SNR Estimation

For automated blind signal detection, an SNR estimate is needed to set the thresholds properly for a particular detection performance. During the conduct of the research, it was found that while arithmetic averaging results in a biased estimate of the input signal strength, geometric averaging produces a reduced bias. As a result, arithmetic averaging tends to estimate the signal-plus-noise energy, while geometric averaging tends to estimate signal energy. This leads to the idea of forming a combination of both estimates to arrive at an SNR estimate. The potential advantage of this method is that traditional SNR estimates rely on identifying regions of the spectrogram that contain noise-only. However, these regions are unknown in many practical scenarios. This new technique may be useful, particularly where persistent and uniform-strength components exist in the input signal. (See also [9], [10].)

5.1.6 Relationship to Other Processes

Literature searches have resulted in the identification of several research areas that appear to be relevant to the current research. These include the concepts of homomorphic processing (see also Appendix B), and myriad filters [11]. These and other processes such as time-frequency transformations (Gabor, Wigner-Ville, etc.) warrant further study.

Appendix A: The Adjustable Bandwidth Concept (ABC) Algorithm

Digital signal processing (DSP) techniques have been developed which use FFT-based power spectral density (PSD) estimates and which average either over time, or over frequency. The ABC algorithm allows for both time and frequency averaging in a unique fashion to be described. It will become apparent that the ABC algorithm is a linear filter process applied to spectral data, and is therefore related to the signal cepstrum.

In the current ABC approach, the spectrogram, a time-frequency representation, is accomplished using the short term FFT. To generate the spectrogram, the input signal to be analyzed is partitioned into N-point (i.e., N-sample) time segments, where N is a power of 2. More specifically, the log-spectrogram is used by the ABC algorithm. The log-spectrogram is formed by logarithmic compression of the magnitude of the coefficients of each of the FFT frequency bins. The time segments can be non-overlapping for computational efficiency, or overlapped for enhanced time resolution. Likewise, various data windows can be used to control spectral leakage. The resulting N-PSD estimates² are input to the ABC process as shown in Figure A-1. As each time-consecutive N-point PSD is input, an N-point output is available at each stage.

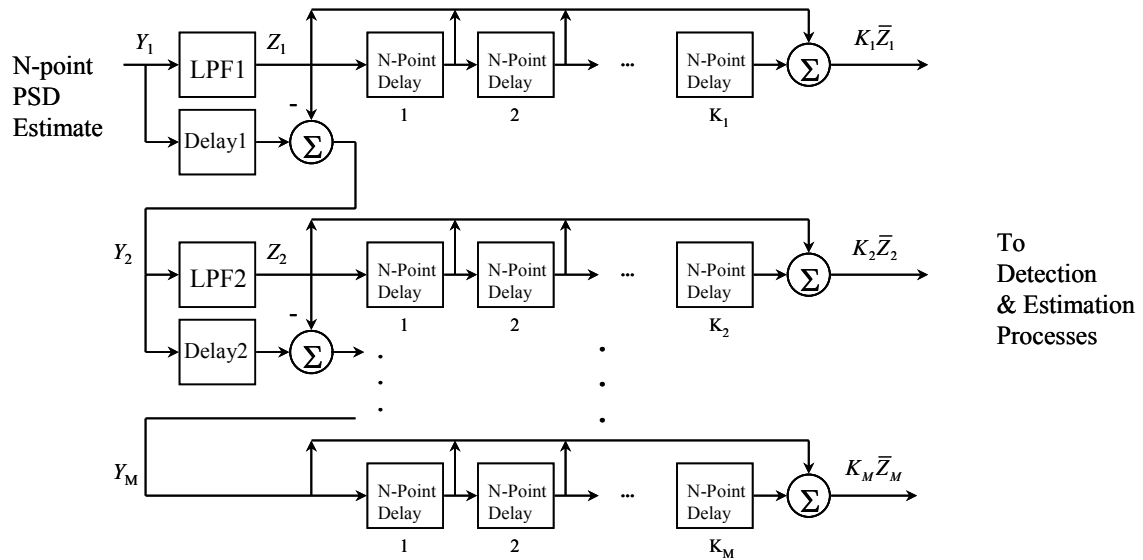


Figure A-1. The ABC algorithm and representative outputs. (U.S. Patent 5,257,211)

Each stage of the M-stage ABC process is similar. In stage M, however, there is no need for averaging over frequency; in this stage only averaging over time is performed as represented by the N-point delays. The summing nodes are implemented as vector sums. A single N-point delay in any given stage yields the moving average of 2

² Just as with a human operator, the performance of the ABC algorithm is directly affected by the quality of the spectrogram. The selection of parameters related to the input spectrogram is an external process to the ABC algorithm.

consecutive PSD estimates; 2 N-point delays yield the moving average of 3 consecutive PSD estimates, and so on. Averaging over frequency is accomplished in the units labeled as LPF1, LPF2, etc. These are symmetric finite impulse response (FIR) filters which are generally low-pass in nature. The number of frequency-consecutive bins that are averaged together is controlled by the order of the filters, which can be up to some fraction of N. For a typical configuration, in stage 1 more averaging over frequency is performed, rather than averaging over time. It is also typical that averaging over frequency decreases with increasing stage number, and averaging over time increases.

A key aspect of the algorithm is the use of unit-gain symmetric FIR filters for effecting the averaging over frequency. Because the data contained in consecutive frequency bins is treated by the filter as if it were a time series, a delay exists on the output samples relative to the input samples. When the filtering is viewed as a smoothing of the frequency data, it becomes apparent that a residual sequence can be generated by vector subtraction of the smoothed sequence from the pre-filtered sequence. However, delays are required to properly align these sequences prior to subtraction. These are shown in the figure as Delay1 through DelayM-1. The residual sequence is passed to the following stage for further processing. (In practice, circular convolution can be performed rather than linear convolution when implementing the filter process.)

In the 3-stage ABC process described in Section 2, $N = 1024$, Delay1 = 20, and Delay2 = 5. Unit-gain Hanning window type FIR filters are used. In stage 1, no averaging over time is performed, while stages 2 and 3 average over 6 and 10 segments respectively. This configuration can be described as allowing for wide-bandwidth detection on the output of stage 1, medium-bandwidth detection on the output of stage 2, and narrow-bandwidth detection on the output of stage 3. It is this bandwidth-based description of the process that the ABC algorithm is named after. The benefit of selecting more or less stages remains an open question. In practice, 3 stages have produced very useful results.

For the results shown in Section 2, a simple threshold detector is used on the output of each stage. The threshold for the first stage should be set based on knowledge or estimation of the noise floor. (For example, the rank-select-threshold (RST) method of noise floor estimation can be used as a method of automated noise floor estimation. In this method, the content of the PSD bins are sorted in ascending strength, and the strength of the bin at around the p^{th} percentile is selected as the noise floor estimate, where p is some appropriate fraction.) Thresholds beyond stage 1 are set independently of the thresholds of all stages. In Section 2, the stage 2 and 3 thresholds were respectively set to 3 and 6 dB. It should be noted that more sophisticated detection methods can be employed. As an example, it may be appropriate to employ hysteresis in the detection process, to avoid detection chatter at signal band edges.

A.1 Averaging Over Frequency

For an M stage ABC process implemented using circular convolution to accomplish the averaging over frequency bins, let $Y_m(k_i)$ represent the input sequence for stage m , with

$$k_i = k + Ni, \quad (\text{A-1})$$

where the indices to the frequency bins are $k = 0, 1, 2, \dots, N-1$. Thus there are N frequency bins available for every time segment, $i, i = 0, 1, 2, \dots$. The input to each stage after stage 1 can be written as

$$Y_{m+1}(k_i) = Y_m(k_i) - Z_m([k_i + d_m]_{\text{mod } N}), \quad m = 1, 2, 3, \dots, M-1. \quad (\text{A-2})$$

Here, d_m is the sample delay due to averaging over frequency in stage m , and is by design, integer valued. Thus, symmetric FIR filters of even order are chosen for averaging over frequency. The input to stage 1 is the log-scaled power spectral density estimate and is represented as $Y_1(k_i)$. Prior to averaging over time segments, the output of each stage is $Z_m([k_i + d_m]_{\text{mod } N})$, for stage $m < M$. For stage M we have no averaging over frequency, such that $d_m = 0$ and

$$Z_M(k_i) = Y_M(k_i). \quad (\text{A-3})$$

The pre-aligned output, $Z_m(k_i)$, for stage $m < M$ is

$$Z_m(k_i) = (1/N) \sum_{l=0}^{N-1} H_m(l) Y_m(l_i) \Omega_N^{-kl}, \quad (\text{A-4})$$

where

$$\Omega_N = e^{-j2\pi/N}, \quad (\text{A-5})$$

and

$$l_i = l + Ni. \quad (\text{A-6})$$

The sequence $Y_m(l_i)$ is related to $Y_m(k_i)$ as

$$Y_m(l_i) = \sum_{k=0}^{N-1} Y_m(k_i) \Omega_N^{kl}. \quad (\text{A-7})$$

The complex-valued response weights, $H_m(l)$, are a result of transformation of the corresponding zero-padded FIR filters, $h_m(k)$, as

$$H_m(l) = \sum_{k=0}^{N-1} h_m(k) \Omega_N^{kl} . \quad (\text{A-8})$$

These filters are normalized such that

$$\sum_{k=0}^{N-1} h_m(k) = 1 , \quad (\text{A-9})$$

and are generally designed such that the filter orders, $2d_m$, have the property

$$2d_{m+1} < 2d_m < N-1 , \quad m = 1, 2, 3, \dots, M-1. \quad (\text{A-10})$$

A.2 Averaging Over Time

When uniform weighting is used, averaging over time is accomplished as the moving average

$$\bar{Z}_m(k + Ni_0 - N \frac{K_m-1}{2}) = \frac{1}{K_m} \sum_{i=i_0}^{i_0+K_m-1} Z_m([k_i + d_m]_{\text{mod } N}) , \quad (\text{A-11})$$

for an arbitrary index offset, $i_0 \geq 0$, and odd number of segments, K_m . The sample delay, $\frac{K_m-1}{2}$, due to averaging over time, is relevant in applications where it is important to synchronize the output to the input. This leads to the reformulation of (A-11) as

$$\bar{Z}_m(k_i) = \frac{1}{K_m} \sum_{i_0=i-\frac{K_m-1}{2}}^{i+\frac{K_m-1}{2}} Z_m([k + Ni_0 + d_m]_{\text{mod } N}) , \quad (\text{A-12})$$

valid at stage m for $i \geq \frac{K_m-1}{2}$ and K_m odd.

Appendix B: Filter Bank Representation of the ABC Process

An alternate representation of the ABC pre-detection process is obtained by considering the equivalent filter bank resulting from frequency averaging. It should be pointed out that because the set of filters comprising the filter bank operates on the log-scaled periodogram, this is sometimes referred to as a lifter bank. The term “lifter” rather than filter is used as a reminder that the lifter bank is defined over periodic components of the periodogram sequence as opposed to periodic components of the time sequence. Likewise, the term quefrency is introduced as a reminder that the periodic components of the periodogram are related to time delay content rather than frequency content of the original input time sequence. See also [12].

It can be shown that the following iterative process is effective for determining the impulse responses, $b_i(n)$, of the equivalent ABC frequency averaging process:

$$\text{Stage 1: } b_1(n) = h_1(n)$$

$$\text{Stage 2: } b_2(n) = h_2(n - d_1) - b_1(n) * h_2(n)$$

$$\text{Stage 3: } b_3(n) = h_3(n - d_1 - d_2) - [b_1(n - d_2) + b_2(n)] * h_3(n)$$

$$\text{Stage 4: } b_4(n) = h_4(n - d_1 - d_2 - d_3) - [b_1(n - d_2 - d_3) + b_2(n - d_3) + b_3(n)] * h_4(n)$$

.

.

.

$$\text{Stage M: } b_M(n) = h_M(n - \sum_{i=1}^{M-1} d_i) - \left[\sum_{i=1}^{M-1} b_i(n - \sum_{j=1}^{M-1-i} d_{i+j}) \right] * h_M(n), \quad (\text{B-1})$$

Here, the filter $h_i(n)$ corresponds to the low-pass filter, LPFi, of stage i. Each symmetric FIR filter, $h_i(n)$, has delay d_i . Also, $h_M(n) = \delta(n)$, and $d_M = 0$. Note that the bracketed term in Equation (B-1) is simply the sum of time-aligned filter bank impulse responses, prior to stage M.

Using this method, Figure B-1 shows the filter bank response for the specific ABC system used in the examples given in Section 2 of this report. It can be noted that with the exception of the Stage 1 filter, the responses are not the typical shapes associated with filter banks. In particular, the responses tend to overlap each other fairly substantially. This in part explains why signals that are present in the spectrogram can often be detected in more than one stage of the ABC output.

Without much difficulty, it is possible to obtain responses that can be considered more typical, as shown in Figure B-2. In this case, the order of the second stage filter is the same as the first stage. A filter order of 64 was chosen for the example, resulting in 65 coefficients for the first and second stage filters. Frequency cutoffs of .05 and 0.2 were chosen for the first and second stages respectively. These cutoffs are relative to a 2 Hz sampling rate. Both filters were Hanning window designs.




Figure B-1. Filter (lifter) bank representation for a specific 3-stage ABC process.




Figure B-2. Filter (lifter) bank representation; Hanning window FIR filter for stages 1 and 2, 65 coefficients and frequency cutoffs of .05 and .2 respectively.

Appendix C: Task Outline

TASK 1: Non-linearity Assessment

OBJECTIVES: To study the effects of non-linearities such as logarithmic compression, on statistics of the spectral estimates. The ABC algorithm is unusual in that averaging is performed after non-linear compression, affecting statistical performance.

APPROACH: Review wealth of open literature work in detection theory to determine the relevancy to the ABC method. Apply well founded methods for analytical development of the underlying statistics of the various ABC output stages.

TASK 2: Theoretical Receiver Operating Characteristic (ROC)

OBJECTIVES: To determine the ROC predicted performance of the ABC algorithm when viewed as a receive system.

APPROACH: The ABC algorithm provides for control over parameters affecting both time and frequency domain averaging. Results from task 1 will be used to generate parameterized ROC performance curves. Particular attention will be needed regarding threshold-based detection. Residual spectral estimates beyond stage 1 of the ABC algorithm should take into account input spectral energy at detected time-frequency locations.

TASK 3: Experimental Receiver Operating Characteristic (ROC)

OBJECTIVES: To verify the ROC predicted performance of the ABC algorithm via experimentation.

APPROACH: An existing MATLAB implementation of the ABC algorithm will be leveraged to help verify predicted performance based on basic input signal scenarios to be determined. Where theory and practice diverge, the theoretical analysis of task 2 will be refined and re-verified experimentally

References

- [1] Rik Pintelon, et. al., "The Geometric Mean of Power (Amplitude) Spectra has a Much Smaller Bias than the Classical Arithmetic (RMS) Averaging," *IEEE Transactions on Instrumentation and Measurement*, Vol. 37, No. 2, June 1988;
- [2] G. Corsini, et. al., "Cramer-Rao Bounds and Estimation of the Parameters of the Gumbel Distribution," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 31, No. 3, July 1995;
- [3] Filippo Attivissimo, et. al., "A Study on Nonlinear Averagings to Perform the Characterization of Power Spectral Density Estimation Algorithms," *IEEE Transactions on Instrumentation and Measurement*, Vol. 49, No. 5, October 2000.
- [4] P. D. Welch, "The Use of Fast Fourier Transform for the Estimation of Power Spectra: A Method Based on Time Averaging over Short, Modified Periodograms," *IEEE Transactions on Audio Electroacoustics*, Vol. AE-15, pp. 17-20, June 1967.
- [5] M. Evans, N. Hastings, B. Peacock, *Statistical Distributions*, Second Edition, New York, J. Wiley, 1993.
- [6] J. G. Proakis, *Digital Communications*, New York, McGraw-Hill, 1983.
- [7] B. Costello, "Graphical Interface Concept for a Signal Detection Process," AFRL-IF-RS-TM-2003-1, February 2003.
- [8] R. Srinivasan, "Robust Radar Detection using Ensemble CFAR Processing," *IEE Proceedings Radar, Sonar and Navigation*, vol. 147, Issue 6, December 2000.
- [9] M. Beauchemin, K. P. B. Thomson, G. Edwards, "The Ratio of the Arithmetic to the Geometric Mean: A First-Order Statistical Test for Multilook SAR Image Homogeneity," *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 34, No. 2, March 1996.
- [10] I. H. Woodhouse, "The Ratio of the Arithmetic Mean to the Geometric Mean: A Cross-Entropy Interpretation," *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 39, No. 1, January 2001.
- [11] J. G. Gonzalez, G. R. Arce, "Weighted Myriad Filters: A Robust Filtering Framework Derived from Alpha-stable Distributions," *IEEE Conference on Acoustics Speech and Signal Processing*, pp. 2833-2836, Vol. 5, May 1996.
- [12] A. V. Oppenheim, R. W. Shafer, *Digital Signal Processing*, Prentice Hall, 1975.